

Philosophy and Theory in Ethics

1.1 Introduction

At the start of his book *Philosophy: A Very Short Introduction*¹ Edward Craig remarks that all of his readers will already, to some extent, be philosophers. That is, he says, they will have commitments about questions of value and about what is the ultimate explanation for the world we inhabit, and the kinds of things we are. If this is what it means to be a philosopher, in however minimal a sense, then clearly we are also, obviously, already ethicists before ever we come to formal study of moral philosophy. Everyone who reads this will have some conception of right and wrong actions, and of the kinds of justifications which may be given for those actions. It is likely that every reader will also have come up against situations where serious disagreements have emerged between people about what ought to be done.

One dramatic example, which is exercising the international community as these words are typed, is whether there should be military action against the political and military leaders of Iraq. None of us, though, is likely to be stuck for personal examples of hard choices and competing convictions.

A distinction which is often made in moral philosophy is between normative ethics (that is, ethical reflection on 'norms', rules on how we ought to behave or limits to the ways in which we may behave), and metaethics, the study of the nature of moral concepts such as goodness and duty, the grounds of moral obligation, and the nature of moral argument. The distinction is a curious one, for all that it does distinguish two tasks of moral philosophy. How can we begin reflecting upon moral norms unless we have some idea of what serves as the basis for those norms? Equally, what good would it be to come up with an impressive piece of philosophical reflection in metaethics which bore no relation whatsoever to any judgements of good and bad action which people actually hold? At best, we might say that morality could be based on such-and-such a view. The fact that no one believes any of the normative commitments which this view implies might mean that everyone's norms are wrong (but how could we find that out?) It might also constitute a basis for impatience with the metaethical view; what point is there in discussing it, if it bears no relation to what people actually believe?

The distinction we have just mentioned is one which became standard through the Twentieth Century. Its legacy in moral philosophy teaching has sometimes been a distinction between 'practical ethics' and 'ethical theory' courses. Apart from any worries we might have about this separation of theory and practice, this has had the unfortunate effect of making courses in practical ethics revolve around a particular implicit moral theory (usually some version of utilitarianism, which we will discuss below), or at least deferring any discussion of theory even at the moments when it was most needed. It also meant that moral theory courses appeared curiously detached from the business of choice and action in the world which

¹ Edward Craig, *Philosophy: A Very Short Introduction* (Oxford: OUP, 2002) p.4

morality seemed to the philosophically unsophisticated eye to be about. A not unnatural impatience frequently set in amongst those studying such courses.

The present course is designed as an introduction to the major areas of moral theory which form the context for current debate. In offering an introduction to theories, I have tried as far as possible to offer practical examples which illustrate some of the implications, and also the limits, of the theories discussed. The danger of this kind of introduction is that the theories can be received as a kind of toolkit, with the appropriate one grabbed from the tool bag when a moral decision has to be made and justified. It is notable that discussions of particular areas of practical ethics (punishment, say) are frequently associated with particular moral theories (in this case, utilitarianism) because it is clear what the theory implies for the relevant area of practical ethics. When we shift focus to, for example, patient autonomy in medical contexts, we find that the theory which gives us the 'right' answer - 'right' means, I suppose, the answer which all suitably reflective liberal citizens of Western democracies believe in - becomes the natural one to discuss. In this case, it will be some sort of Kantian theory which emphasizes issues of justice and autonomy.² Yet each of the theories mentioned was meant to be global descriptions of the nature of morality and moral thinking, intended to encompass all of our practical ethics, not only the bits where they gave the answers we already believed anyway.

In doing moral philosophy, we are engaged in a process which seeks to give a consistent account of the human good, interpreted in terms of action (what we should do) and character (what we should be), along with the issue of why it is that these things are what we should do and be. If the content of this course has a bias towards theory, it is because this is the area of moral philosophy which most people have not reflected on systematically; it is certainly not intended to imply that moral theory can be detached from questions of moral practice.

Much moral philosophy over the last two centuries has centred upon questions of action, of what ought and ought not to be done. Where moral philosophy has become integral to the concerns of disciplines which see their most pressing moral task as answering questions of conduct - as in a good deal of contemporary medicine, military strategy, political policy - this is natural enough. In the fifth chapter we will refer in a more extended fashion to the revival of interest in the ethics of character. This area of character is in fact relevant to all of the disciplines just mentioned, medicine, the military and so forth, as well as offering a context for reflection upon the kind of question which moral philosophy asks and seeks to answer when it enquires 'how can we be human, well?' That is, how can we be free from the variety of frustrations which prevent us from living well? The frustrations which might spring most immediately to mind are illness and poverty, but others, like lack of human company, or of opportunities for learning and for play, could well figure here too. Placing the question 'what is it to be human, well?' at the heart of ethics makes an important assumption, that there are better and worse ways of being human, and that it is central to the business of ethics to discern these.

² Medicine provides an intriguing test case here. When discussing the ethics of health care distribution, the moral theory invoked is often overtly utilitarian; when discussing the relationship of doctor and patient, in terms of duty to tell the truth and responsibility not to violate trust, the justification for the duty may well be anti-utilitarian. One of the central challenges to contemporary medical ethics is the creation of a coherent moral vision; those who believe that our moral visions may be true or false will hope that, as well as being coherent, the moral vision developed will also be true!

1.2 'Morals' and 'Ethics'

We must make an explicit note regarding terminology. The roots of the two word-groups, on the one hand ethic, ethics, ethical, etc., and on the other moral, morals, morality belong within different languages - classical Greek, and Latin. Each, though, concerns pretty much the same territory: the area of custom. The emphasis which these words give to the notion of the customary is slightly different. Our word 'moral' comes from a Latin root which has a strong sense of social expectation. 'Ethics' is rooted in a Greek term which carries more of a sense of individual character, or what is customary for an individual, what he or she is disposed to do. Whilst etymology gives no clue to current meaning (and certainly doesn't prescribe any corrective to it), it is worth emphasizing that no really useful distinction between the two groups of terms looks likely to arise from their history. The fact that modern dictionaries of the English language frequently define each group in terms of the other may encourage us to think that there is no fine distinction to be drawn here. There is a temptation to assume that when we meet phrases such as 'moral ethics' we are simply confronting tautology (or perhaps confusion), and perhaps we should give in to that temptation.

Having said this much, though, we ought to note that philosophers have sometimes found it useful to mark out certain distinctions within this area, and they have found the ready availability of two word-groups a handy one for labelling the things that they wish to distinguish. Thus, Bernard Williams distinguishes between 'ethics' (which he does not define very closely, but which concerns the recognition of a range of phenomena such as obligations, duties and so forth), and 'the system morality', by which he means a more systematic approach, which he says

should be understood as a particular development of the ethical, one that has a special significance in modern Western culture. It peculiarly emphasizes certain ethical notions rather than others, developing in particular a special notion of obligation, and it has some peculiar presuppositions.³

We do not need to press the question of what Williams sees as the content of 'morality' in order to recognize that he is marking out a distinction in which 'ethics' is, for him, a vaguer object of study, but none the worse and perhaps much the better for that. Philosophers may have played their part in the development of 'the system morality' but it could not have gained a 'special significance in Western culture' if it were solely a philosophical development.

In this section we will offer a brief account of the more and less ambitious ways in which philosophy may have a role to play in inquiring about the good, and goodness. It is proper to note at the outset, though, that these central terms, 'ethics' and 'morality', may be used in quite specific and distinct ways by philosophers, marking out distinctions which they wish to

³ Bernard Williams, *Ethics and the Limits of Philosophy* (Cambridge, Massachusetts: Harvard UP, 1985), p6. The thought that there is an important distinction to be marked around the concept of obligation as it occurs in modern moral philosophy does not originate with Williams. G.E.M. Anscombe in her article 'Modern Moral Philosophy' (available in a number of collections, for example as essay 4 in her *Collected Philosophical Papers Volume III* (Oxford: Basil Blackwell, 1981)), suggests that the concept of obligation had a role only where there was still a pervasive belief in a focus for that notion - one who obligates, which in terms of a Western civilization shaped by Christianity will be labelled 'God'.

make but which do not map on to any equivalent distinction in everyday language. There seems no strong reason why a philosopher, wanting to make the same distinction Williams makes, should not have used the labels in the opposite way, with 'ethics' as the more systematic and 'morality' as the less. Williams's use is non-arbitrary in the sense that the distinction is supported in a marginal way by classical usage of the words, but arguably modern usage would make 'ethics' the system and 'morals' (which has a stronger sense of plurality) the less formal object of study.

1.3 What is the Question?

We will not worry too much, then, whether our study is termed 'moral philosophy' or 'philosophical ethics'. What, though, is it a study of or about? A typical answer might well be that moral philosophy is about what should be done - about what we ought to do, or what we have a duty to do (or ought not to do, or have a duty not to do). Interestingly, though that does indeed seem to many people to be the very first question in ethics, it has not always seemed so. To the classical world, it seemed rather that 'what sort of life should I lead?', or 'how can I live well?', was the more basic question. 'What is the good life, the life lived well?' appears a much broader question than the very specific 'what ought I to do?', and its very breadth may daunt us. We may in any event have no faith at all that any sort of answer could be sustained to a question like 'what is the good life?' 'What ought I to do?', when it is asked deliberatively, in a particular situation and by a particular person, seems to offer a greater likelihood of being subject to reasoned reflection towards a conclusion.

The idea that we might start with questions of action and not of character, of events in a life rather than of the nature of a life, is part of the move which we saw Williams identify with 'the system morality'. Impressed by the evaluation of acts as the first point of departure for ethical thought, moral philosophy is pressed into offering theoretical accounts of what actions are to be done (or not to be done), and why (or why not). We will look at some of these accounts through the sections of this module. It is important, though, to see what is at stake in such moral theories. The first thing which is at stake we have already noticed; it is a picture of what assessment of the moral life amounts to, namely an attachment of an evaluative claim to an action or a potential action. The second thing is the criteria for success or failure which we bring to the theory. There are, obviously enough, criteria by which any theory in any discipline might be regarded as a bad theory. Inconsistency would be one; a theory which has some internal commitments which end up claiming that $0=1$ would be a bad theory. A theory which does not 'fit the facts' would also be a bad theory. What would it be, though, for a moral theory to 'fit the facts'?

The third chapter will confront the question of the place of facts in ethics more directly. All we need do for the moment is to notice that a moral theory may fail to measure up to moral judgements which we assume to be right. (Perhaps we just find their being wrong inconceivable: most people, when faced with the claim 'genocide is bad', will not make a first response of the form 'well, it all depends...') One *motivation* for developing theories may be to tell us what to do in new and perplexing situations. Human cloning might be one of those. A moral theory might issue in a judgement about what should be done. That same theory might, though, imply a radical reassessment of some *existing* permission or prohibition, one which we are reluctant to believe we could be mistaken in holding. Is this

enough to discredit the theory? Compatibility with our existing moral intuitions might, after all, be one of the 'facts' we want the theory to fit.

We might think of moral theory as playing a dual role for us. It should give a theoretical account of the content of morality as already believed in; and it should offer answers to puzzling questions about what ought to be done in new situations. If we think of moral theories in this way, we may expect them to be consistent with the content of our existing moral beliefs, and regard it as a failure of the theory if it is not. Braver theoreticians, though, suggest that the failure of their theories to match our normal expectations of the content of morality might be a measure of the need for revision in our everyday moral expectations. Uncomfortably aware that there seem to have been significant changes in moral outlook in the past, concerning slavery, the role of women in society, and so forth, situations where we have come to believe that we were wrong, we might think that being challenged over our normal moral expectations is not in the end such a bad thing. Was it ever theory, though, which caused a change of mind in these cases?

One of the most significant and influential contemporary moral theories is, as it happens, one which conflicts in a variety of ways with many people's ordinary moral intuitions. It is a moral theory (or perhaps a cluster of theories) which centre on the consequences of action as the basis for moral evaluation. The most discussed of these 'consequentialist' theories is that termed 'utilitarianism'. In its simplest form, utilitarianism proposes that the effects of an action in promoting or inhibiting pleasure and the absence of pain will serve as a basis for moral judgement on that action. In comparing the moral value of actions, we may say that the more pleasure produced, the better the action; and in deliberating over how to act, we should ask what action (including, to anticipate later discussion, doing nothing at all) would produce the greatest amount of pleasure. It is not our concern in this immediate context to say more about this theory than simply to note that a variety of actions which might normally be regarded as prohibited by ordinary morality could plausibly be required by utilitarianism.

To take a simple example: Henry has just had a blood test and is waiting to hear if he is HIV positive. His doctor, Joan, believes that it would be better from a utilitarian perspective in Henry believed he was HIV positive; knowing his character, she suspects that if he were to believe this then some aspects of his lifestyle would change, both to his own benefit and to that of his friends and colleagues. He is not, in fact, HIV positive. Should she tell him that he is?

A utilitarian might say that she should lie, for Henry's own good and that of other people. On the other hand, a utilitarian might say that she should tell him the truth. The utilitarian might justify this by pointing out the disutility which would result if it were to become known that Joan did not tell the truth. Her patients might regard her diagnoses and advice as merely well-meaning opinions regarding their welfare, perhaps unconnected with their actual bodily condition. This is something which would itself have consequences, very plausibly bad ones, worse ones than telling the truth to Henry could have. What matters to us, though, is not which of the judgements about consequences is correct. Any outrage which 'ordinary morality' would feel at Joan's failure to tell the truth to Henry would relate precisely to its being a failure to tell the truth, in a situation where telling the truth is expected and required. What ordinary morality sees as a (simple) violation of duty, the utilitarian explains in different terms, as an incorrect calculation of consequences. If the consequences had stacked up differently, it could be right to lie.

The strength of this complaint against utilitarianism, that it fails to capture the judgements of

ordinary morality, rest in this case on a sense, which many people share, that there are certain duties which cannot be violated, whatever the consequences. To this thought the utilitarian may reply by wondering out loud whether this sense of duty ‘whatever the consequences’ is one which we could honestly be willing to commit ourselves to on sober reflection. We can all imagine circumstances where doing what we take to be our duty could have utterly devastating consequences. (Try this: if I do not tell a lie, the villain will press the button and the nuclear bomb will explode.) Thankfully, we are rarely in situations where this question has any practical force for us, and it may simply be wrong to imagine that once the utilitarian has exposed a willingness on our part to (say) tell a lie to save a friend’s life, we are thereby committed to the utilitarian cause ourselves. Maybe we should give an account of the variety of duties which we believe ourselves to have, and a complex account of the ways in which under certain conditions certain duties may trump others. Such an exercise would hardly be a trivial one, and the nature of duty (the reason why these things are duties in this and not that circumstance; the reason why this duty trumps that one in this circumstance) would not even have begun to be broached. Compared to this, utilitarianism has one great advantage, that of economy. Its key elements can be stated in just a few lines. We have seen one of its apparent disadvantages, namely that it fails to conform to ordinary morality. This is not yet a decisive argument against it, and it has many advocates, especially in practical ethics.

1.4 Some key terms and ideas

No academic discipline is devoid of jargon, and moral philosophy has plenty of its own. A good dictionary of philosophy is a useful tool; even professional philosophers may make use of them from time to time! We conclude this first chapter with a few terms used to classify moral theories, some of which will be discussed further below. Some readers may find it helpful to use this section for reference as necessary while studying the rest of the module than to attempt to read and digest it all at one go.

Theories which seek to determine right action by reference to the consequences of our actions are often lumped together under the heading ‘consequentialist’. We have already met this rather ugly word, and noted the most common consequentialist moral theory, utilitarianism. This theory will concern us in the next chapter. To the consequence ‘pleasure and the absence of pain’ which forms the basis of utilitarianism, we might offer as alternatives ‘the extension of the realm of human choice’, ‘the maximization of Gross Domestic Product’, ‘the expansion of the range of human knowledge’, and perhaps any number of others. Moral theories labelled ‘consequentialist’ are often contrasted with those termed ‘deontological’ (that is, based in a conception of duty), but, as we shall see when we discuss Kant in chapter 4, consequentialists often have a very highly developed sense of duty. There are undoubtedly contrasts to be drawn between those theories based on consequences and those which are not, but the contrast is not best located in duty.⁴

Another contrast which is often offered is between *teleological* and *non-teleological* moral theories. (The phrases come from the Greek word *telos*, meaning ‘end’ or ‘goal’.) Theories labelled teleological may contain a vision of what good human living would look like, and

⁴ See chapter 4 section 2.

justify their substantive or normative moral content by showing how it contributes towards the achievement of the goal. Marxist political theories have something of this form, as do many capitalist ones; indeed, the ultimate goals of the two are often stated in quite similar terms, leading to the question whether the two are simply alternative means to a similar end! Utilitarian moral theories are often labelled teleological, but are unlikely to have any clear conception of the ultimate form of good human living. Indeed, beyond a vague sense of 'maximal pleasure' some noted utilitarians have explicitly disavowed any more contentful picture of where their moral theory might at best lead. Kant's conception of ethics includes quite a well developed vision of the fulfilment of the moral life, but he denies that this can serve any role in particular moral actions or decisions, so that his 'telos' is detached from his account of the practical details of the moral life. It is therefore unclear whether to count him as a teleological or a non-teleological moral philosopher; it will depend on the criteria we adopt for 'teleological', and in spite of the confidence with which the word is often used there is no general agreement on how it should be used which would decide this matter for us.

A further area of distinction is between *realist* and *anti-realist* conceptions of morality. Moral realists will assert that there are moral truths which are true independently of any human beliefs about them; in theory, at least, all human beliefs about morality could be false, and something (God's will, maybe, or a special realm of metaphysical 'moral facts', or something else) would be what makes it the case that the beliefs are false. There will then be an important debate about the means of access which human beings might have to this realm of 'moral reality'. Maybe the answer will begin in a theological category like revelation, or in an account of metaphysical intuition (a sort of parallel with sight in the physical world). Maybe it will be in nature, avoiding the apparently mysterious areas of religion and metaphysics altogether. Those who are skeptical about the grounding of morality in natural facts (see chapter 3 below), and who are unimpressed by the idea of justifying a moral claim through revelation or metaphysical intuition, might well come to the view that there is actually nothing which could justify such claims; they are not "grounded" anywhere, in any reality. Moral anti-realism, then, is likely to amount either to a denial of the status of truth or falsity to any moral claim, or possibly to some new 'non-realist' account of what the truth or falsehood of a moral claim consists in.

One direction which such anti-realist thought might take is despair over the whole business of ethics; another is to see morality as a human cultural construct, with moral claims being claims made relative only to the culture which they occur within. This would be a form of *relativism*, in which (at least in this form) any truth which a moral claim such as 'incest is wrong' might have is a truth within a culture. There would be no 'universal' ethics. Or at least, there would not be any universal ethics unless, by accident, all cultures happened to agree on what is right and wrong. We might imagine the development of a kind of world-culture which had a fairly stable, shared moral outlook; such a thing might well be happening in business practices to a considerable degree as global trade develops. We should notice, though, that this 'global moral culture' would not have developed as a response to a moral reality, or to the objective existence of a moral realm independent of human perception. To this extent, moral relativism and anti-realism are compatible with everyone having the same moral views, just as moral realism is consistent with everyone having different moral views (and all of them perhaps being false!) To *relativism* is often contrasted *objectivism*, and to each of these a third category, *subjectivism*.

The moral subjectivist holds that the use of moral evaluative terms like 'good' and 'bad' has its origins and justification not in an objective realm of moral facts, nor in a human culture,

but in certain facts about the human subject, such as desires, opinions, feelings or attitudes. There are a number of good treatments of some aspects of subjectivism, and I do not extend the discussion here but refer the reader to the essay suggestions at the end of the chapter. It is worth noting, though, that one moral theory, that of Kant, has features which may make us wonder whether it should be counted as subjectivist or objectivist. As we discuss Kant's theory in chapter 4 it may be worth bearing this question of classification in mind, not least as an instructive exercise in the limits of usefulness of some of the central terms of classificatory art.

Our final terms are *cognitivist* and *non-cognitivist*. A cognitivist theory will be one which has room for a concept of moral 'cognition' - coming to know that a certain moral claim is true or false. As suggested earlier, this might include a very wide range of candidates, such as 'knowledge through religious revelation', as well as a belief in moral facts accessible through the senses. Moral non-cognitivism will be applied to the variety of moral beliefs in which there are no moral facts or truths to be known. Typically, non-cognitivists will feel a need to give some account of the kind of things moral claims are. For example, if I say 'abortion is wicked' I might be expressing an attitude, or prescribing a form of practice which I enjoin on others. Neither the expression of the attitude nor the prescription are claims of truth. One reason to be resistant to the simple identification of moral cognitivism with moral realism is that if there is room for a contextualized form of moral truth in some forms of non-realist or anti-realist theory (as I take the form of relativism sketched above to be), there might also be forms of contextual cognition which determine that the contextual moral truths are truths in that context, though this is not a theme to be pursued here.⁵

This sketch of some 'terms of art' has been undertaken in part to assist the reader when approaching other works of moral philosophy, but also in part to signal some of the distinctions which one might have in mind when approaching the study of moral philosophy. To a moral realist, it might be appropriate to speak of the work of Human Fertilization and Embryology Authority as attempting to 'discover' the truth of various moral questions; a non-realist might be more inclined to speak of inventing answers to novel questions, claiming that there is no pre-existing 'truth of the matter' at all. But "realism" and 'non-realism' classify theories, and are not themselves theories. It is to the first of our discussions of moral theories that we now turn.

⁵ Another reason to resist the identification is that a realist might hold that there is an objective basis for moral truth which is nonetheless absolutely inaccessible to human sense and reason, so beyond cognition. (As a parallel, imagine the possibility of a new, sixth sense which gives access to a dimension of reality which is actually 'out there' but which we do not know about.) Why anyone would hold this to be true is another question. Some theological ethics appears to have a flavour of this approach about it, for example that of Karl Barth.